


Hamilton-Jacobi-Bellman (HJB) Equation (Justin Ling)

In these notes, I will be using a lot of approximations that you should be comfortable with.

- ① $f(x_0 + dx) \approx f(x_0) + \partial_x f dx$ (Euler Approximation for small dx) ($\partial_x f$ is partial derivative with respect to x)
- ② $\int_0^{dx} f(\tilde{x}) d\tilde{x} \approx f(0) dx$ (Leibniz Rule for small dx)

To understand the HJB equation, we need to understand what a Value Function is, and why it's important. A Value Function is defined as such:

③ $V(x_0, t_0) :=$  $\min_{u[t_0, t_f]} \left\{ \int_{t_0}^{t_f} \mathcal{L}(x(t), u(t)) dt + \mathcal{P}(x(t_f)) \right\}$

- $x \in \mathbb{R}^n$ is the state, and $x_0 = x(t_0)$ is the initial state.
- $u \in \mathbb{R}^m$ is the control, and $u[t_0, t_f]$ is a function defined on the interval $[t_0, t_f]$, so $u[t_0, t_f] := \{ u(t) \mid t \in [t_0, t_f] \}$
- t_0 and t_f are initial and final times. Note that t_f is fixed, while t_0 is variable.
- $\mathcal{L}(\cdot, \cdot)$ is a running cost. Integrating over $\mathcal{L}(\cdot, \cdot)$ gives us a cost over a trajectory.
- $\mathcal{P}(x(t_f))$ is a terminal cost. Later we will see what happens when $\mathcal{P}(\cdot) = 0$.

④ $\dot{x} = f(x, u)$ is the time derivative of x .

Oftentimes we know how to penalize each point in time with $\mathcal{L}(x(t), u(t))$, but trying to minimize the integral $\int_{t_0}^{t_f} \mathcal{L}(x(t), u(t)) dt$ requires a search over trajectories $u[t_0, t_f]$. However, searching over trajectories is difficult. If possible, we should try to simplify so that we ~~can~~ can evaluate a favorable $u(t)$ at only one point in time, using an evaluation of favorable states x .

Complex 😞

Favorable Trajectories
 $u[t_0, t_f]$
 $x[t_0, t_f]$



Simple 😊

Favorable state x
Favorable control $u(t)$ for one point in time

There are ways to do this outside the scope of this lecture.

To evaluate favorable states x , we'd like a simple map $x \mapsto \mathbb{R}$ that weights the states, with higher weights being unfavorable. You can think of this like an energy function.

This is exactly what our Value Function (3) is doing, with a small twist. It is a function of both state and time, not just state. Therefore, the Value Function is a constant landscape function with respect to a fixed interval $[t_0, t_f]$.

The Value Function reveals an interesting property if we split the integral by a small dt .

$$\textcircled{5} \min_{u[t_0, t_f]} \left\{ \int_{t_0}^{t_0+dt} \mathcal{L}(x(t), u(t)) dt + \int_{t_0+dt}^{t_f} \mathcal{L}(x(t), u(t)) dt + \mathcal{P}(x(t_f)) \right\}$$

$$\stackrel{\text{why } \textcircled{2}}{\approx} \min_{u[t_0, t_f]} \left\{ \mathcal{L}(x_0, u(t_0)) dt + \int_{t_0+dt}^{t_f} \mathcal{L}(x(t), u(t)) dt + \mathcal{P}(x(t_f)) \right\}$$

First, we define $u_0 := u(t_0)$. Notice the expression $\mathcal{L}$ on the left $\mathcal{L}(x_0, u_0)$ only depends on u_0 , so the expression can be rewritten as,

$$\min_{u_0} \left\{ \mathcal{L}(x_0, u_0) dt \right\} + \min_{u[t_0+dt, t_f]} \left\{ \int_{t_0+dt}^{t_f} \mathcal{L}(x(t), u(t)) dt + \mathcal{P}(x(t_f)) \right\}$$

$$\underbrace{\hspace{10em}}_{\min_{u_0} \left\{ V(x(t_0+dt), t_0+dt) \right\}} \textcircled{6}$$

The right expression is also a Value Function! It's identical to the original (3), except incremented forward by small dt . The integral is now only dependent on the interval $[t_0+dt, t_f] \Rightarrow [t_0, t_f]$, but it is implicitly dependent on u_0 through the initial condition $x(t_0+dt) \approx x_0 + f(x_0, u_0) dt$. Therefore, we only need to minimize (6) over u_0 , and the minimization over $u[t_0, t_f]$ happens through the definition of $V(x(t_0+dt), t_0+dt)$. Now we have,

$$V(x_0, t_0) = \min_{u_0} \left\{ \mathcal{L}(x_0, u_0) dt + V(x(t_0+dt), t_0+dt) \right\}$$

By the Euler Approx. (1), $V(x(t_0+dt), t_0+dt) = V(x(t_0) + f(x_0, u_0) dt, t_0+dt) = V(x_0, t_0) + \frac{\partial V}{\partial x}(x_0, u_0) f(x_0, u_0) dt + \frac{\partial V}{\partial t}(x_0, u_0) dt$

$\frac{\partial V}{\partial x}$ is a gradient

Using this approximation, dropping the (x_0, u_0) dependencies in V for readability,

$$V(x_0, t_0) = \min_{u_0} \left\{ \int_{t_0}^{t_f} L(x, u) dt + \partial_x V^T f(x, u) dt + \partial_t V dt \right\} + V(x_0, t_0)$$

Cancel $V(x_0, t_0)$ on both sides

↓
not dependent on u_0 , so move outside min-brackets

$$\frac{d}{dt} \left[\partial_t V + \min_{u_0} \left\{ L(x, u_0) + \partial_x V^T f(x, u_0) \right\} \right] = 0$$

↳ must be 0,

$$\Rightarrow -\partial_t V = \min_{u_0} \left\{ L(x, u_0) + \partial_x V^T f(x, u_0) \right\} \leftarrow \text{HJB Equation.}$$

Now we have an equation to relate x_0 and u_0 , instead of $u[t_0, t_f]$, $x[t_0, t_f]$.

Note on terminal cost $P(x(t_f))$.

Note that $V(x_0, t_f) = P(x(t_f))$, since the integral collapses. So the shape of $P(x)$ has an impact on the system's behavior near t_f . This is important in cases like MPC where you expect your system to continue acting past t_f .

We can derive a "near- t_f " HJB with the similar methods as before.

$$\begin{aligned} V(x_0, t_f - dt) &= V(x_0, t_f) - \partial_t V dt = \min_{u[t_f-dt, t_f]} \left\{ \int_{t_f-dt}^{t_f} L(x(t), u(t)) dt + P(x(t_f)) \right\} \\ &\quad \text{using (1)} \quad \text{using (3)} \\ &= \min_{u[t_f-dt, t_f]} \left\{ L(x_0, u_0) dt + P(x_0 + f(x_0, u_0) dt) \right\} = P(x_0) + \min_{u_0} \left\{ L(x_0, u_0) + \partial_x P^T f(x_0, u_0) \right\} dt \end{aligned}$$

Since $V(x_0, t_f) = P(x_0)$, cancelling on both sides gives you the HJB equation. For the min condition to be satisfied, we take the gradient of $L(x_0, u_0) + \partial_x P^T f(x_0, u_0)$ and set it to 0.

$$\partial_{u_0} L(x_0, u_0) + \partial_x P(x_0)^T \partial_{u_0} f(x_0, u_0) = 0$$

Here, we can see that the optimal value of u_0 , which represents the value of $u(t)$ for any t near t_f , the terminal cost induces a significant offset and tradeoff. The term $\partial_x P^T f(x_0, u_0)$ emphasizes the relevance of x_f implicitly. This is much easier to see in the discrete-time case.

~~In the discrete case, it is easy to see that without terminal cost, u_0 is not affected by the terminal cost. It depends on x_0 .~~